

Trust and Partner Selection in Social Networks: An Experimentally Grounded Model

Riccardo Boero*, Giangiacomo Bravo[†] and Flaminio Squazzoni[‡]

Abstract

This paper presents an experimentally grounded model on the relevance of partner selection for the emergence of trust and cooperation among individuals. By combining experimental evidence and network simulation, our model investigates the link of interaction outcome and social structure formation and shows that dynamic networks lead to positive outcomes when cooperators have the capability of creating more links and isolating free-riders. By emphasizing the self-reinforcing dynamics of interaction outcome and structure formation, our results cast the argument about the relevance of interaction continuity for cooperation in new light and provide insights to guide the design of new lab experiments.

Keywords: trust; investment game; cooperation; laboratory experiment; agent-based model; network structure; partner selection.

1 Introduction

Recent experimental studies have explicitly acknowledged the relevance of trust for cooperation in human societies (e.g. Barrera & Buskens 2009; Berg *et al.* 1995; Boero *et al.* 2009a,b; Camerer 2003; Keser 2003). From a sociological perspective, the problem with these studies is that they took into account stylized and highly unrealistic interaction structures, e.g., random coupled subjects, so that the structure of the game was nothing but a sequence of one-shot interactions. This was due to the standard laboratory approach of economic studies that is aimed to pursue the potential effect of the embeddedness of agents in complex social structures. On the contrary, many formal models with a more concrete sociological content investigated exactly how structural embeddedness affected trust and cooperation, but without

*Dipartimento di Scienze Economiche e Finanziarie “G. Prato”, Università di Torino and GECS — Research Group in Experimental and Computational Sociology

[†]Dipartimento di Scienze Sociali, Università di Torino, Collegio Carlo Alberto and GECS — Research Group in Experimental and Computational Sociology. Corresponding author. Mail address: Dipartimento di Studi Sociali, Via S. Ottavio 50, 10124 Torino, Italy. Email: giangiacomo.bravo@unito.it

[‡]Dipartimento di Studi Sociali, Università di Brescia and GECS — Research Group in Experimental and Computational Sociology

referring to empirically verified assumptions on agents' behavior (e.g. Cohen *et al.* 2001; Pujol *et al.* 2005).

This paper aims to study the link between interaction outcomes and social structure formation by combining experimental evidence and social network simulation through an agent-based model (from now ABM). Unlike influential studies that emphasized the relevance of interaction continuity for the emergence of cooperation (e.g. Axelrod *et al.* 2002; Cohen *et al.* 2001), we investigate whether the robustness of cooperation depends on the capability of agents to select their partners and whether this might be influenced by the particular network structure where they are embedded in. For this purpose, the advantage of combining laboratory experiments and simulation models is twofold. On one hand, experiments provide sound and clean data on agents' behavior on which to build informative micro-macro models. Unfortunately, this is a not so common practice, since most cooperation models have been based on stylized theoretical assumptions rather than on specific empirical evidence so far. On the other hand, ABM can extend experimental evidence by providing for the impossibility of exploring complex interaction structures and long-term macro dynamics and evolution in the laboratory. In this mutual benefit, there is a cross-fertilization that could be largely beneficial for sociological research (e.g. Boero & Squazzoni 2005; Chmura & Pitz 2007; Dal Forno & Merlone 2004; Ebenhöf & Pahl-Wostl 2008; Janssen & Ostrom 2006; Rauhut & Junker 2009; Rouchier 2003).

In light of this inspiration, we have extended the scope of a standard experimental economics game, the repeated investment game, to situations that were impossible to explore in the laboratory by introducing the effects of complex network structures. Our main substantial result is that different static network structures do not lead to outcomes substantially different from random networks. A substantial increase in cooperation occurs only when partner selection and dynamic structure formation are introduced, which allow cooperators to break their links with free-riders and let them to benefit from an higher number of interactions.

Our results show that the robustness of cooperation among agents, rather than being favored by the stability of the interaction structure, as suggested by previous studies (e.g. Axelrod *et al.* 2002), have more to do with the selective intelligence of agents when they give rise to networks based on people who like each other that isolate bad apples. Moreover, our experimentally-grounded results confirm the findings of Eguíluz *et al.* (2005) model, where the cooperative equilibrium in a Prisoner's Dilemma played on a dynamic network was guaranteed by "leaders" (agents in central position) that played a crucial role in sustaining cooperation in the system. In fact, also in our case, the presence of a cooperative agent in a central position in dynamic networks had a crucial influence in shaping a structure favorable to cooperation composed of trustful and trustworthy agents.

The remaining of the paper is organized as follows. The next section presents the research background. Section 3 describes the experiment used to gather the data. Section 4 presents the estimation of agents' parameters. Section 5 introduces the models and presents the results. Finally, Section 6 discusses our findings.

2 Research background

Most studies on cooperation assumed that individuals have no choice over which opponents they play (e.g. Axelrod 1984, 1997). This implied that, especially in large populations, cheaters may escape punishment by roving from partner to partner, thereby exploiting co-operators. The simulation literature has suggested that the stability and continuity of the interaction structure might provide a solution to combat exploitation by cheaters in social interaction. Simulating an iterative prisoner's dilemma, Cohen *et al.* (2001) have suggested that even just the continuity of the interaction patterns can make cooperative regimes come true in hostile environments, since the "shadow of the future" established through it preserved the contexts in which cooperation strategies worked, allowing later variants and emulations to find a supportive environment. In light of this, Axelrod *et al.* (2002) have emphasized that the persistence of interaction patterns was enough to guarantee cooperation, thereby dispensing from more complex linkage pattern mechanisms.

Although important, these findings did not capture the full picture of cooperation in human societies and underestimated more powerful social mechanisms that can sustain it. A prominent social mechanism that permits cooperators to decrease the risk of being cheated in many real social interactions, by isolating cheaters, is partner selection. It is likely that in concrete social situations individuals have preferential choice and can opt out from given interactions (e.g. Slonim & Garbarino 2008). Partner selection is not just a mean for individuals to do good preferential choices, but it also provides an incentive to the counterparts to be reliable and committed to others so as to avoid social isolation (e.g. Ashlock *et al.* 1996). For example, Joyce *et al.* (2006) have built a simulation model largely inspired by Axelrod's tournament that demonstrated that a conditional association strategy, i.e., assuming that agents left partners who defected against them and stayed with cooperative partners, was better performing than the conditional altruism postulated by Axelrod with the Tit-for-Tat strategy. A similar finding on the relevance of contingent strategies capable of discriminating good and bad partners was also found in a simulation by Aktipis (2004).

These interactional aspects have been explained by the strategic uncertainty reduction and the "commitment bias" arguments, with the former emphasizing rational strategic aspects and the latter intrinsic other-regarding or emotional motivations. The first position has been stressed by Kollock (1994) in an experimental study on market transactions, where the challenge to deal with information asymmetries and uncertainty about the others' trustworthiness made long-term interaction partners look more attractive than others. Podolny (2001) suggested that investment bankers that operated in markets characterized by a higher degree of uncertainty were likely to interact with colleagues they have interacted with in the past. Gulati (1995) found the same evidence in an empirical study on firm alliances. More recently, Beckman *et al.* (2004) have confirmed this in an empirical study on interlock and alliance networks for the 300 largest US firms during the 1988-1993 period, where they found that there was a strong correlation between the type of market uncertainty and the stable/unstable nature of partner selection networks: firms facing firm-specific uncertainty were less selective and more open to new ties than firms facing collective market uncertainty. Hauk (2001) developed a simulation model of an iterated prisoner's dilemma where partner selection was a

learning strategy for agents in uncertain situations, with agents that learn how to discriminate between trustworthy and untrustworthy partners and to reward/punish others as well. By revising strategies according to selected partners, cooperators learn how to avoid that cheaters' exploitative strategies thrive in the population. In two experiments conducted in the US and Japan, Yagamashi *et al.* (1998) have found that uncertainty promoted commitment formation between partners but, in particular, that individuals trusting more others tended to form committed relations less frequently than skeptics, since they relied more on generalized trust.

The second position argued that the risk and uncertainty of exchange provided the opportunity for partners to demonstrate their trustworthiness so that behavioral commitments and trust signals can sustain cooperation in an informal and spontaneous way, without negotiations, bargaining and binding agreements as in most social exchange literature (e.g. Molm *et al.* 2000). Recent simulation studies have suggested that partner selection can generate a "commitment bias", so that people's commitment to existing partners can increase beyond instrumental reasons, thanks to the strengths of intrinsic other-regarding motivations (Back & Flache 2006, 2008). Back (2010) has confirmed this hypothesis creating a lab commitment-dilemma game based on a market exchange between subjects who played in artists selling paintings' roles and computer-simulated collectors buying them, in six locations of three countries, i.e., the Netherlands, China and the US. The findings have showed that the initial commitment of buyers had a positive effect on partner selection, even when controlling for uncertainty reduction and material benefits.

Although full of insights of great importance for sociological analyses of social interaction, these studies did not take into the due account the relevance of complex social structures, and in particular the consequences of partner selection within certain social structure configurations. Most of the experimental literature on partner selection have typically investigated the internal aspects of the selection decision, rather than the consequences of it over time for the emergence of trust and cooperation (e.g. Haruvy *et al.* 2006; Kagel & Roth 2000). Recently, Slonim & Garbarino (2008) have investigated the relevance of partner selection for trust and altruism in the lab, by playing an investment game and a variant of the dictator game, where subjects' partners were exogenously imposed or endogenously chosen. Nevertheless, they underestimated the relevance of studying partner selection sociologically within network formation processes, that is by investigating how network structure and partner selection can influence each other. In fact, partner selection is to be viewed as a driving force for network formation, stability and change (Beckman *et al.* 2004).

At the same time, the abundant literature on network formation that explained important aspects of cooperation in social and economic situations was more analytically oriented than explicitly based on experimental evidence, did not take into account large social systems (e.g. Calvó-Armengol 2001; Dutta *et al.* 2005; Jackson & Watts 2002; Jackson & Wolinsky 1996) and did not investigate the crucial role of partner selection as a driving force for the network dynamics (e.g. De Vos *et al.* 2001). An interesting exception was Eguíluz *et al.* (2005), who elaborated a social network simulation based on a spatial prisoner's dilemma to investigate the co-determination of individual strategies and network, but without partner selection at the agent level and lab data for the modeling of agent decisions.

Our paper aims to fill this gap by combining experimental evidence and network sim-




Figure 1: The investment game: player A profit = $e - I + R$; player B profit = $e + 3I - R$.

ulation to study partner selection, trust and network configurations. This is to understand relevant interaction aspects that is impossible to look at in the laboratory and to provide new insights to improve the sociological content of current economic laboratory experiments.

3 The experiment

Our goal was to observe individuals' behavior in asymmetric trust situations, i.e., where a trustor has to choose whether placing trust or not in another individual (the trustee) who has a rational incentive to cheat (Coleman 1990, 177–180). The experimental economics standard for studying these situations was the investment game, a protocol based on Berg *et al.* (1995). The game has been repeated many times using different samples and under a vast range of different conditions (e.g Cronk 2007; Keser 2003; King-Casas *et al.* 2005; Knoch *et al.* 2009; Ortmann *et al.* 2000) and, recently, it has gained interest also among sociologists (e.g. Barrera & Buskens 2009; Boero *et al.* 2009b; Buskens & Raub 2008; Molm *et al.* 2000).

Figure 1 shows a schematic representation of a typical investment game. Participants were coupled and randomly assigned to different roles, called player A (the trustor) and player B (the trustee). Both players received an endowment of e ECU. First, player A had to decide the amount I between 0 and e to send to B, keeping for him/herself the part $(e - I)$. The amount sent by A was tripled by the experimenter and sent to the trustee, in addition to his/her own endowment, who decided how much to return to A. As before, the amount returned by B could be any integer between 0 and $(d + 3I)$. The amount $(d + 3I - R)$ not returned represented B's profit, while R was summed to the part kept by A to form the final profit of the latter, which was hence calculated as $(e - I + R)$. All profits earned during the game were counted in experimental currency units (ECU), translated in real money using a fixed exchange rate and paid immediately after the end of the experiment. The game is called “investment game” since the rule of tripling the amount sent by the trustor implies that (a) the trustor deals with the uncertainty of paying a cost at the beginning of the interaction to gain possibly higher revenues at the end, and (b) the trustee has a return from the trustor's decision.

Since B had no rational incentive to return anything to A, the dominant strategy for A players was to invest zero. This led to a unique subgame perfect equilibrium, where both players kept their entire endowments. However, this was inefficient since any sum grater

than zero invested by A was tripled by the experimenters and could lead to a Pareto superior outcome, with the optimum represented by A investing the whole endowment. Despite these equilibrium predictions, a quite robust result of investment games played both in the laboratory and in the field was that A players invested a substantial amount of their endowments (usually around 40%) and that B players returned slightly less than the amount invested by A , although much less than a fair share of the overall amount ($d + 3I$) that they received (e.g. Berg *et al.* 1995; Boero *et al.* 2009b; Cronk 2007; Keser 2003; Ortmann *et al.* 2000).

Our experiment was based on a repeated version of the investment game as described above. One hundred and eight subjects were recruited through public announcements and played the game in six groups of 18 individuals. Half of the subjects were students of the University of Brescia, half of the University of Torino (Cuneo campus). The experiment took place in Spring 2009 in the computer laboratories of the Faculties of Economics of the two universities, which were both equipped with the experimental software z-Tree (Fischbacher 2007). All groups received exactly the same instructions and played a repeated investment game using an identical computer interface. The endowment was 10 ECU with an exchange rate of 1 ECU = 2.5 Euro. Participants earned, on average, approximatively 15 Euros, including a show up fee of 5 Euros, that were paid immediately after the experiment. The experiment, including instruction reading, took approximatively one hour.

All interactions took place through the computer network and the subjects were unable to identifying their counterparts. Participants played a repeated investment game for a total of 10 periods and were informed in advance of the number of periods to play. The experiment used a “stranger” matching protocol, i.e., randomly re-coupling of subjects after each period. This protocol permitted to control for the use of direct reciprocity strategies among participants. The players’ roles were randomly assigned at the beginning of the first period and subsequently changed on a regular basis, so that each subject played the same number of period in both roles.

The results show that subjects invested, on average, 3.48 ECU and returned 2.79 ECU, slightly less than what is usual in similar games.¹ This can be explained if subject pool used, which was composed of undergraduate students in business and economics, is taken into account. This population is well known for under-average cooperation in experiments and this is even more true when subjects played investment and trust games (see Camerer 2003, 65).

4 Parameter estimation

We used these experimental data to calibrate an ABM that reproduced the behavior of the subjects. With regards to A players, we estimated a coefficient β_i that indicated how much the player modified his/her investment each round in function of the difference between the amount invested and the amount received by B players in the previous round. For any player i and round t , we computed the difference $X_{it} = R_i - I_i$, where I_i and R_i were the amounts that i invested and received as return from his/her investment in the previous round respectively.

¹All statistical analysis were conducted using the R 2.10.1 platform (R Development Core Team 2009).




Figure 2: Distribution of the estimated coefficients

We subsequently fitted the model

$$Y_{it} = \alpha_i + \beta_i X_{it} + \varepsilon \quad (1)$$

where Y_{it} was the amount invested by player i in round t , in order to obtain two parameters α_i and β_i for each subject that defined his/her behavior as A player. The equation (1) took into account that A players could have had an individual constant propensity to trust represented by the individual intercept α_i , but also the capability of reacting upon past experiences, which was captured in the β_i coefficient.

On the other hand, B players were supposed to react mainly against what they received from A players. In order to capture their behavior, we estimated a third coefficient γ_i defined as the average amount returned by each subject as proportion of the amount received in each round plus the fixed endowment. Therefore, the parameter γ_i represented an estimate of the player's trustworthiness.

We were able to successful estimate the parameter for 105 out of the 108 subject that participated in the experiment. Figure 2 presents the distribution of the estimated coefficients. It is worth noting that, while β_i did not significantly correlated with the other parameters, the correlation between α_i and γ_i was significant and positive ($r = 0.44$, $p < 0.001$). In other terms, trustful subjects tended to be trustworthy as well: a result consistent with the hypothesis that agents who have a general propension toward cooperation behave similarly in different situations..

5 The models

5.1 Experiment replication

We first built an *experimentLike* model that replicated the original experiment.² In each period agents were coupled and played an investment game either as *A* or as *B* player. At the end of the periods they were randomly re-coupled while their roles alternated on a regular basis. The behavior of each agent depended on the coefficient estimated from the experimental data. More specifically, each agent i playing in the role of *A* in period t invested an amount

$$I(i, t) = \alpha_i + \beta_i [R(i, t - 1) - I(i, t - 1)] \quad (2)$$

where α_i and β_i were the coefficients of the corresponding participant while $I(i, t - 1)$ and $R(i, t - 1)$ were the amounts invested and received in the previous round. In the first round of the game, we set $I(i, 0) = R(i, 0) = 0$ that made agents choosing their first move on the basis of α_i alone. This was justified since α_i reflected the intrinsic propensity of the participants (and, henceforth, of our agents) to trust others players, while β_i represented the effect of past experience. On the other hand, when playing on the *B* positions agents returned

$$R(i, t) = \gamma_i [3I(i, t) + 10] \quad (3)$$

which was simply their endowment multiplied by the individual parameter γ_i estimated on the experimental data.

Based on this setting, we built an *experimentLike* model, which succeeded in reproducing the experimental data.³ The average investment over 100 run of the *experimentLike* model was 3.57 ECU and the average return 2.76 ECU (see Tab. 2). A t test over individual investments/returns, averaged over the 10 periods of the game, confirmed that these figures did not significantly differ from the experiment (respectively, $t = 0.288$, $p = 0.774$ two sided, and $t = -0.085$, $p = 0.933$ two sided).

5.2 The simulation scenarios

Once checked the model capability of reproducing the experimental data, we designed new scenarios to test the plausible effects when important variables in the experimental design were modified. All scenarios used the experimental data as input and, more generally, were identical to *experimentLike* model except for the specific characteristic(s) under examination. Table 1 provides an overview of all simulation scenarios.

A first simple variation of the game design was to increase the number of periods played by agents. This has been done for all scenarios, including the *experimentLike* model, that

²Note that we drawn from the data 105 valid estimations out of 108 subjects. Being the game played in dyads, we decided to remove randomly one of the experiment participants at the beginning of each run of the model, which, as a result, included only 104 agents.

³All models were programmed in Java using the JAS Library (<http://jaslibrary.sourceforge.net/>). Interested readers can obtain the source code upon request by the authors.

Model name	Main characteristics
<i>experimentLike</i>	<ul style="list-style-type: none"> • random coupling in each period • one way interaction
<i>twoWays</i>	<ul style="list-style-type: none"> • random coupling in each period • two way interaction
<i>fixedCouples</i>	<ul style="list-style-type: none"> • fixed couples • two way interaction
<i>denseNetwork</i>	<ul style="list-style-type: none"> • fixed fully connected network • two way interaction
<i>smallWorld</i>	<ul style="list-style-type: none"> • fixed small-world network • two way interaction
<i>scaleFree</i>	<ul style="list-style-type: none"> • fixed scale-free network • two way interaction
<i>dynamic1Couples</i>	<ul style="list-style-type: none"> • dynamic network • broken links are replaced only for isolated agents • two way interaction • start from random coupling
<i>dynamic1Dense</i>	<ul style="list-style-type: none"> • dynamic network • broken links are replaced only for isolated agents • two way interaction • start from dense network
<i>dynamic2Couples</i>	<ul style="list-style-type: none"> • dynamic network • broken links are replaced only by one of the two formerly linked agents • two way interaction • start from random coupling
<i>dynamic2k10</i>	<ul style="list-style-type: none"> • dynamic network • broken links are replaced only by one of the two formerly linked agents • two way interaction • start from from a regular network of degree 10

Table 1: The simulation scenarios

were tested both using a 10 and a 30 period game. Then we explored the introduction of two way interaction among agents, i.e. in each period agents played both as *A* and *B* with their opponent (*twoWays* model). Note that, while this extension of the model could be of no particular interest in itself, it provided a baseline for subsequent scenarios which extended the interaction to a number of agents greater than two.

Since we knew that differences in the network structure can have crucial consequences

for the cooperation at a macro level, we built four models that explored the effect of different network structures. The *fixedCouples* model simply maintained the initial couples throughout the whole run. The *denseNetwork* model introduced a fully connected network where, in each period, every agent interacted with all other agents. In the *smallWorld* model, we introduced a fixed “small-world” network structure (Watts & Strogatz 1998; Watts 1999) with initial degree 4 and re-wiring probability 0.01. The *scaleFree* model was based on a “scale-free” network that conformed to Barabási & Albert (1999) algorithm with a degree equal to 3. In all these models, the network structure remained fixed throughout the whole game. Note that agents may have more than one interaction per period. More specifically, they played an investment game both as *A* and as *B* with all other agents with whom they were linked in the current period.

Finally, two models were introduced to investigate the effects of the introduction of a dynamic network in the game. Basically, these scenarios added a criterion of partner selection to the formation of couples, so that interaction structure was dynamically shaped by the outcome of the interaction (e.g. Corten & Buskens 2009; Eguíluz *et al.* 2005; Flache 2001). Since there was no way to derive from the experiment the algorithm that agents used to break their links and create new ones, we identified certain plausible algorithms using as few assumptions as possible. We built two different models, both based on simple mechanisms but with a slightly different focus. The first one (named *dynamic1*) was designed to avoid that any agent became isolated — i.e. no longer linked to any other agent — while the second (*dynamic2*) was aimed at maintaining constant the total number of links in the network.

Both models were based on the idea that “unsatisfied” agents could break their links. In other words, we assumed that agents had a threshold happiness function so that they changed partners when they were unhappy. The threshold function was very simple: *A* agents were happy when *B* returns were higher or equal to the ones in the previous time step. This makes sense on the light of some fundamental attitudes of human beings as highlighted by economic psychology (e.g. Kahneman & Tversky 2000). Note also that a similar strategy for determining agents’ satisfaction has been assumed, for instance, in Bravo (2010) work that modeled the emergence of institutions in common-pool-resource situations. We assumed that a couple could be broken when at least one agent of the couple was unhappy. When there was no past information to evaluate the happiness — i.e. in the first two periods — the couple remained fixed.

The only difference between the two dynamic models regarded the algorithm for the definition of the new links. In the *dynamic1* model, broken links were replaced only when an agent became isolated. In this case a new link was formed between the isolated agent and a new agent different from the one with whom it was linked before. In the *dynamic2* model, after each link break, one of the two formerly connected agents was chosen at random to initiate a new link. In this case, no check was made and agents may become isolated. This second algorithm maintained constant the number of links in the network, while changing its structure at the same time.

Both the *dynamic1* and the *dynamic2* model were tested starting from different initial network structures. More specifically, we used random coupling and a fully connected network as starting point for the former and random coupling and a random network of average de-

Model name	10 period game		30 period game	
	A investments	B returns	A investments	B returns
<i>experimentLike</i>	3.57 (2.50)	2.76 (2.62)	3.56 (2.53)	2.76 (2.63)
<i>twoWays</i>	3.57 (2.52)	2.76 (2.61)	3.57 (2.54)	2.76 (2.61)
<i>fixedCouples</i>	3.65 (2.53)	2.91 (3.13)	3.67 (2.56)	2.92 (3.17)
<i>denseNetwork</i>	3.57 (2.54)	2.76 (2.61)	3.57 (2.54)	2.76 (2.61)
<i>smallWorld</i>	3.58 (2.54)	2.76 (2.62)	3.57 (2.54)	2.76 (2.62)
<i>scaleFree</i>	3.61 (2.54)	2.80 (2.68)	3.61 (2.54)	2.80 (2.69)
Experiment	3.48 (2.69)	2.79 (3.58)	–	–

Table 2: Average investments and returns in the original experiment and in the simulation scenarios. Standard deviations are in parenthesis. Simulation results significantly different from the experimental ones are highlighted using a bold font.

gree 10 for the latter. The difference was due to the fact that the second re-coupling algorithm determined a situation where each agent was connected to every other and there was no other possibility of re-linking than re-building the broken connection. Having two link replacement algorithms and two initial network structure, we tested four dynamic models named, respectively, *dynamic1Couples*, *dynamic1Dense*, *dynamic2Couples* and *dynamic2k10*. All dynamic models were run for 30 periods in order to see the evolution of the network in a time longer than the original experiment.

5.3 Static network model results

In order to take into account the stochastic elements of the models, we ran 100 replications for each scenario. Table 2 presents the average investments and returns in all models, comparing them with the experiment.

The first important point is that neither the modification of the network structure nor the increase of the number of periods changed significantly the outcome of the game. In most of the runs, models reached their equilibrium before the 10th period, maintaining it subsequently with no major change. Even the introduction of two-way interactions, which increased the number of investment games played by agents in every run, simply accelerated the approaching to the equilibrium without altering the results. Similarly, letting agents play over a fully connected network, where each agent interacted with all other agents in every single period (resulting in a total of 214,240 interactions per run in the 10 period game and in 642,720 interactions in the 30 period one) did not change the overall results. In all these cases, according to a *t* test, the average investments and returns of the models did not significantly differ from the experiment.

The introduction of networks where not all the agents had an equal number of interactions posed a challenge for our analysis. Up to now, to check the statistical significance of the difference between the simulation results and the experiment, we used a *t* test over individual averages. This worked fine as long as all the agents interact an equal number of times, while

Model name	α_i	β_i	γ_i
<i>experimentLike</i>	-0.75***	0.08	-0.91***
<i>twoWays</i>	-0.74***	0.10	-0.92***
<i>fixedCouples</i>	-0.71***	0.22*	-0.85***
<i>denseNetwork</i>	-0.73***	0.12	-0.92***
<i>smallWorld</i>	-0.72***	0.11	-0.91***
<i>scaleFree</i>	-0.46***	0.14	-0.69***

Table 3: Correlation between average payoffs per run and agents' parameters. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

it lead to an overestimation of the weight of agents with a lower number of interactions in case of networks where the number of links for each agent were different. For instance, the average investment in the 10 period *scaleFree* model was 3.61 when considering each single interaction. However, taking in consideration the average contribution for each agent, the resulting average was 3.57. Although small, this difference could affect the significance of our results. Moreover, it was likely that the difference increased moving from static to dynamic networks. In order to take into account this issue, we performed our t tests on the weighted means, where the weights were obviously given by the number of interactions performed by each agent.

The *smallWorld* and *scaleFree* model results were very similar to the previous ones: the average amount invested were slightly higher than the experiment, but a t test showed that this difference was not significant. Also the average returns were almost identical to the experiment.

An interesting question was whether the change in the network structure would have changed the relative advantage of cooperators (i.e. trustful and trustworthy individuals) over free-riders. We estimated the relation between the value of the three parameter defining each agent and the payoff earned in the different scenarios. Since average investments and returns did not change moving from a 10 to a 30 period game, we took into account only the 10 period game. Table 3 reports the correlation coefficient between the agent parameters and its earnings.

In most models, earnings significantly and negatively correlated with both α_i and γ_i , while the correlation with β_i was not significant. In other words, both trustful and trustworthy individuals earned lower payoffs in all games, while the agent ability to react to others' action did little difference. There was one noticeable exception to this rule, namely the *fixedCouples* model where the correlation coefficient relative to β_i , although small, was positive and significant at the 5% level. This means that reactive agents, by increasing their investments after a high return and decreasing them following a low return, were somewhat advantaged in this setting. This was probably due to the fact that the fixed couple structure of the game tended to favor reciprocity based strategies. Another partial departing from the rule mentioned above was the *scaleFree* model. In this case the relation between the agents' trust (measured by α_i) and their earnings was much weaker than in any other situation. To some extent, the same

Model name	Period 1–10		Period 11–20		Period 21–30	
	A invest.	B returns	A invest.	B returns	A invest.	B returns
<i>dynamic1Couples</i>	3.65 (2.58)	2.92 (2.96)	3.67 (2.60)	2.95 (2.90)	3.68 (2.62)	2.96 (2.93)
<i>dynamic1Dense</i>	3.79 (2.67)	3.32 (3.20)	3.66 (2.60)	2.96 (2.96)	3.68 (2.62)	2.97 (2.94)
<i>dynamic2Couples</i>	3.82 (2.68)	3.37 (3.42)	4.48 (3.01)	5.02 (4.50)	4.63 (3.11)	5.58 (5.12)
<i>dynamic2k10</i>	4.11 (2.82)	4.00 (3.59)	4.43 (3.01)	4.85 (4.30)	4.49 (3.04)	5.02 (4.50)
Experiment	3.48 (2.69)	2.79 (3.58)	–	–	–	–

Table 4: Average investments and returns in the original experiment and in the dynamic network models. Standard deviations are in parenthesis. Averages significantly different (at the 10% level) from the experimental ones are marked in bold.

was true for their trustworthiness (γ_i).

5.4 Dynamic network model results

Moving from a static to a dynamic network led to noticeable results (see Tab. 4). Note that, being the network no longer static, the amount sent and returned in the different parts of the game varied. Therefore, we separately analyzed the results of periods 1–10, 11–20 and 21–30 of the game. The average cooperation in the *dynamic1* models slightly increased, even if a t test (performed on the weighted means for the reasons given above) showed that the difference was significant only for returns in the first ten rounds of the *dynamic1Dense* model ($t = 1.446$, $p = 0.075$, one sided). The *dynamic1Dense* model resulted in higher cooperation levels than the *dynamic1Couples* one because cooperative agents, starting from 103 links each, were initially able to sever most of them, thereby avoiding be exploited by free-riders. However, this strategy became no longer possible once the situation settled down to an average of 1–2 links per agent, with the consequence that trustful agents’ earnings decreased below those of the less trustful agents.

Especially interesting was the network dynamic. Apparently the system had a single equilibrium, which was reached independently from the starting point. From period 9 onwards, the average degree of the networks resulting from the *dynamic1Couples* and the *dynamic1Dense* models converged to the value of 1.67 ± 0.02 (Fig. 3, left panel). This was true although the initial degrees of the two networks were, respectively, 1 and 103. The final network structure of the two models was also very similar with, in both cases, all agents linked in small groups of 2–8 individuals (Fig. 3, center and right panels).

Most of the times, the center of the group was represented by an especially cooperative agent, a fact consistent with the findings of Eguíluz *et al.* (2005) model, where agents played a Prisoner’s Dilemma on a dynamic network and the equilibrium was guaranteed by “leaders” (agents in central positions) that played an essential role in sustaining cooperation in the system. In the network plotted in the middle of Figure 3, agents with more than two links had an average α_i of 4.02, while the one for agents with two links or less was 3.75. Similarly, agents with more than two links had an average β_i of -0.21 (vs. 0.22) and an average




Figure 3: The system resulting from the *dynamic1* models converged to a fixed equilibrium independently from the starting point. The left panel shows the average degree of the networks resulting from the two models. The center and the right panel report the networks resulting after 30 periods of a typical run of the *dynamic1Couples* and the *dynamic1Dense* model respectively.

γ_i of 0.21 (vs. 0.13). More generally, both α_i and γ_i of agents were positively correlated with the number of interactions in the last 10 periods of the game. In the *dynamic1couples* the correlation coefficients were, respectively, 0.42 and 0.86, while in the *dynamic1dense* they were 0.40 and 0.78.⁴ In other words, agents with more links (thereby performing more interactions) were both more trustful and more trustworthy.

As for the static models, we studied the relation between the agents' parameters and their earnings. Table 5 report our estimations. The results of the *dynamic1* models did not depart much from the static network ones. Again trustful and trustworthy agents were disadvantaged, while the capacity to react made little difference. One noticeable exception appeared in the final periods of the *dynamic1Couples* model, where the coefficient related to the β_i parameter became significant. This was rather surprising, because the small stable emerging groups (Fig. 3, center panel) were not very different from the dyads that characterized the *fixedCouples* model, which also produced a significant correlation between β_i and agents' earnings.

Another exception was found in the initial periods of the *dynamic1Dense* model, where the correlation coefficient for α_i was almost zero and the one for γ_i was positive and highly significant. However, this situation did not persist and in the period 11–20 the coefficient became already positive and highly significant. The initial advantage for trustworthy agents was due to their capacity of maintaining their links longer, while untrustworthy lost theirs because of the unsatisfaction of their opponents. Similarly, trustful agents were not worse off than the others because they quickly broke their unsatisfying links. However, once the network settled up, the situation changed and the coefficients became similar to the *dynamic1Couples* model.

The *dynamic2Couples* model led instead to a strong increase in cooperation (Tab. 4). Returns were already significantly higher than in the experiment in the first 10 periods of the game and further increased subsequently. Investments were significantly higher than in the

⁴All estimates were significant at the 1% level.

Model name	Period	α_i	β_i	γ_i
<i>dynamic1Couples</i>	1–10	−0.78***	0.11	−0.80***
	11–20	−0.78***	0.16	−0.67***
	21–30	−0.68***	0.21*	−0.66***
<i>dynamic1Dense</i>	1–10	0.02	−0.18	0.64***
	11–20	−0.72***	0.17	−0.66***
	21–30	−0.66***	0.15	−0.61***
<i>dynamic2Couples</i>	1–10	−0.13	−0.05	0.53***
	11–20	0.37***	−0.14	0.96***
	21–30	0.37***	−0.14	0.95***
<i>dynamic2k10</i>	1–10	0.32***	−0.19	0.87***
	11–20	0.36***	−0.18	0.96***
	21–30	0.35***	−0.19	0.97***

Table 5: Correlation between average payoffs per run and agents’ parameters. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

experiment from the 10th period onwards. Starting from the same period, returns were higher than investments as well. In other words, investments became profitable for *A* players. This was also clear when looking at the correlations presented in Table 5. The correlation between the agents’ payoffs and both their α_i and γ_i became *positive* and significant, i.e., trustful and trustworthy agents earned, on average, higher payoffs. The explanation is that more trustful and trustworthy agents had more links, so enjoying a higher number of exchanges than less trustful/trustworthy ones. More precisely, the correlations between α_i and γ_i of each agent and the number of interactions that the same agent performed during the whole 30 period game were both positive and significant ($r = 0.38$, $p < 0.001$ and $r = 0.92$, $p < 0.001$ respectively). Vice versa, less cooperative agents became easily isolated and earned, on average, lower payoffs.

The resulting network was also remarkably different from the *dynamic1* model. Many agents were now isolated, while a large cluster emerged around the most cooperative agents (Fig. 4). Therefore, the most trustful/trustworthy agents were the most connected. For instance, the average α_i of agents with zero or one link in Fig. 4 was 3.41, while the average α_i of agents with two or more links was 4.57. Similarly, the average γ_i of agents with zero or one link was 0.10, while the average γ_i of agents with two or more links was 0.22.

The *dynamic2k10* model presented a strong increase of cooperation (Tab. 4). Unlike before, both the average investments and returns were higher than in the experiment already in the first ten periods. Investments and returns further increased in the subsequent periods, even if they finally settled down at a slightly lower value than in the *dynamic2Couples* model. As before, investments became profitable for *A* player from period 10 onwards. The correlations between agents’ parameters and payoffs followed a path similar to the *dynamic2Couples* model, with the exception that the correlation between the agents’ α_i and their payoffs was




Figure 4: Networks resulting after 30 periods of a typical run of the *dynamic2Couples* model (left) and of the *dynamic2k10* model (right).

positive and significant from the beginning of the game (Tab. 5). Overall, trustful and trustworthy agents enjoyed a larger number of interactions. The correlation between agents' parameter and number of interactions was significant at the 0.1% level for both α_i and γ_i (while it was negative and significant only at the 10% level for β_i).

The network resulting from a typical run of the *dynamic2k10* model was presented in Figure 4 (right). Unlike the previous case, no agent was now isolated. However, cooperative agents maintained a larger number of links. For instance, the average α_i of agents with five or less links was 2.73, while for agents with more than five links this figure is 4.17. Similarly the average γ_i of the former group was 0.07, while the one for the latter group is 0.17.

6 Discussion

In this paper, we have presented an experimentally grounded ABM that extended the results of lab experiments on trust to study the impact of the interaction structure on the cooperative outcomes. The first important lesson that we can draw from our results is that it is not the stable or unstable nature of the interaction structure *per se* which has a relevant impact on trust diffusion and cooperation. Our result supported the hypothesis that what really matters is the link of interaction outcome and structure formation. While a simple change of the network structure did not alter significantly the outcome of our game, the introduction of a dynamic network led to higher cooperation when cooperative agents had the capability to create more links and, possibly, to isolate free-riders.

This first result stands in stark contrast with the popular literature emphasizing the effects of network structure on interaction outcomes (e.g. Barabási 2009; Buchanan 2002; Watts 1999). The “shadow of the future” postulated by Axelrod (1984) and Cohen *et al.* (2001) certainly represents an important factor in explaining cooperative outcomes. However, once trust has been eroded in static networks by untrustworthy actions, the stability of links between agents using conditional strategies can only produce a never ending run of mutual defection. On the contrary, by permitting agents to leave their partners cooperation increases

since the exploitation opportunities of free-riders are reduced. This is exactly what happened in our simulations under the *dynamic2* conditions. Trustful and trustworthy agents were no longer bound to reciprocate defection, but were in position to find new partners with whom they started cooperative relations. As a consequence, overall cooperation increased, bringing higher fairness and efficiency in the system.

This has implications for sociological and experimental analyses on cooperation and social norms. Previous simulation models highlighted the relevance of tags for cooperation (e.g. Axelrod *et al.* 2004; Hales 2000). In sociological terms, tags can be viewed as communication signals that circulate among people and connote certain attributes of individuals or groups. For instance, they include speaking dialects or particular accents, dressing elegant formal get-ups, or driving garish customized sport-cars. Although largely unintentional and often inaccurate or even false, these attributes or actions might be socially used as predictors of the others' behavior, in particular when interacting with unknown partners, although they can be unintentional, inaccurate or false. As suggested by the simulation model presented in Riolo *et al.* (2001), the presence of tags in a social system can help cooperation to emerge even without past experience, memory, reciprocity motives or reputational signals. This makes partner selection different than in our case, where it was based on past experience of each agent individually. The point is that, if individuals can detect tags or exchange some communication signals before interacting, it is likely that the intelligibility of partner selection and the consequent structure formation can be strongly influenced. This is where the dark side of partner selection is likely to enter the picture, i.e., when tag-based selection is biased by discriminative social beliefs and gives rise to inefficient network formation by preventing unknown partners from potential interaction benefits. If this is reasonable, future laboratory experiments that compare the impact of social tags and past experience on partner selection and structure formation should be particularly welcomed.

Another point is that, if it is reasonable to expect that interaction outcomes matter to give rise to social structures, as we have stressed, it should be essential to investigate empirically and experimentally how social expectations about the trustworthiness of potential partners and the trusting-other attitudes of individuals are influenced by cultural aspects, institutions and normative values at a social level. This is of particular interest in the study on the nature and evolution of markets, as suggested by North (2005). Firstly, particular market institutional scaffolds might be more efficient than others because they have institutionalized partner selection as a relatively low-cost sanctioning for cheaters. Secondly, particular market structure configurations might have historically evolved out of partner selection mechanisms, as suggested also by White (2004), so that complex market structures could be nothing but a magnification of relatively simple social mechanisms..

A second lesson concerns precisely the relevance of partner selection. The increase of cooperation via partner selection has implication both for the efficiency of the system and for its fairness. In the *dynamic2* conditions the average investment increased by one third and the average return almost doubled in comparison with both the experiment and the static networks. While the standard in the experimental literature on trust and cooperation often assumed randomly re-matching couples, as though the game were a sum of one-shot periods, with no chance for partner selection, our results suggested that cooperation can crystallize

in stable interaction structures that were likely to depend on the positive outcomes generated by agent interaction. This means that it is reasonable to suppose that individual action and interaction structure tend to co-evolve in the social life (e.g. Buskens & Weesie 2000; Buskens *et al.* 2008; Mark 1998). Our experimentally grounded model allows us to emphasize that one of these self-reinforcing co-evolutionary links could be driven by partner selection (e.g. Corten & Buskens 2009; Flache 2001; Eguíluz *et al.* 2005). If this is reasonable, by excluding from the radar screen of the experimental analysis these aspects, it is likely that something crucial for explaining social mechanisms of cooperation gets lost.

The third lesson concerns the research methods. By combining laboratory experiments and simulation models, we were able to exploit some of the advantages of both. This is not a common practice, even if earlier works that followed our same inspiration have already suggested some of its potentialities (e.g. Boero *et al.* 2010; Dal Forno & Merlone 2004; Deadman *et al.* 2000; Jager & Janssen 2003; Rouchier 2003). Being empirically defined, the distribution of strategies used by our agents was more plausible than what was commonplace in research based on formal models. At the same time, thanks to simulation, we were able to extend the experimental results beyond the limits of the laboratory (e.g., a larger number of interactions, more complex network structures, etc.).

The link between experiments and simulation is not limited to the creation of empirically-grounded models. Verifying in the lab hypotheses deriving from simulation research is becoming relatively common. A recent example of experimentally tested ABM is Centola *et al.* (2005) “Emperor’s Dilemma”, whose conclusions have been tested by Willer *et al.* (2009). Our simulation evidence can therefore guide the design of a new set of both field studies and laboratory experiments explicitly targeted to study partner selection behavior or, more generally, to gain a better understanding of the relation between individual preferences and the dynamics of social networks. For instance, our model supports the hypothesis that allowing individuals to select their interaction partners increases *per se* the amount of both trust and trustworthiness in a given system. This could be checked in field studies by comparing contexts of social interaction that vary in their degree of individual freedom in choosing their partners (e.g., firm and consumer behavior in monopolistic vs. competitive markets). It is also reasonable to suppose that partner selection can cover relevant functions of informal social control and sanctioning in many social and economic situations when the partner choice and the break of existing links becomes a socially shared information. This opens the door to interesting investigations on the relations between partner selection and reputation both for social order, as a means for norm abiders to isolate and punish cheaters (e.g. Pinyol *et al.* 2008), and for environmental exploration, as a means to know what is the best partner choice in uncertain situations (e.g. Boero *et al.* 2010). The same hypothesis could be checked with more accuracy in the lab by designing a trust-game experiment where individuals can decide whether maintaining or not their current relations, even if this comes at the cost of reducing the number of interactions and the range of potential partners.

To sum up, by linking trust and partner selection our results emphasize new research directions. What is especially important is that the level of analysis moves from static to dynamic networks. Real world networks are rarely fixed and independent of individual actions. They evolve over time as the result of a myriad of independent choices done by agents em-

bedded therein. Taking into account the aggregate level effects of these choices is complex and will plausibly involve a great deal of both empirical and computational new research. However, this will nicely improve our knowledge of fundamental social mechanisms and, *a fortiori*, of important social situations.

Acknowledgements Financial support was provided by a FIRB 2003 grant (Strategic Program on Human, Economic and Social Sciences) from the Italian Ministry for the University and the Scientific Research (SOCRATE Research Project, Protocol: RBNE03Y338_002) and by the Regione Piemonte, Research Project ICT4Law. A preliminary version of this paper has been presented at the Sixth European Social Simulation Conference, held in 2009 in Guilford, Surrey, UK. We thank three anonymous conference reviewers and the audience for useful remarks. The authors gratefully acknowledge comments and suggestions by Jakob Grazzini, Cristina Mosso and Károly Takács.

References

- Aktipis, A. 2004. Know When to Walk Away: Contingent Movement and the Evolution of Cooperation. *Journal of Theoretical Biology*, **231**, 249–260.
- Ashlock, D., Smucker, M. D., Stanley, E. A., & Tesfatsion, L. 1996. Preferential Partner Selection in an Evolutionary Study of Prisoner’s Dilemma. *BioSystems*, **37**, 99–125.
- Axelrod, R., Riolo, R., & Cohen, M. D. 2002. Beyond Geography: Cooperation with Persistent Links in the Absence of Clustered Neighborhood. *Personality and Social Psychology Review*, **6**(4), 341–346.
- Axelrod, R., Hammond, R. A., & Grafen, A. 2004. Altruism via Kin Selection: Strategies that Rely on Arbitrary Tags with which They Evolve. *Evolution*, **58**(8), 1833–1838.
- Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, Robert. 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton: Princeton University Press.
- Back, I. H. 2010. Commitment Bias: Mistaken Partner Selection or Ancient Wisdom? *Evolution and Human Behavior*, **31**, 22–28.
- Back, I. H., & Flache, A. 2006. The Viability of Cooperation Based on Interpersonal Commitment. *Journal of Artificial Societies and Social Simulation*, **9**(1), 12.
- Back, I. H., & Flache, A. 2008. The Adaptive Rationality of Interpersonal Commitment. *Rationality and Society*, **20**(1), 65–83.
- Barabási, Albert-László, & Albert, Réka. 1999. Emergence of Scaling in Random Networks. *Science*, **289**, 509–512.
- Barabási, Albert-László. 2009. Scale-Free Networks: A Decade and Beyond. *Science*, **325**, 412–413.

- Barrera, Davide, & Buskens, Vincent. 2009. Third-Party Effects on Trust in an Embedded Investment Game. *In: Cook, K., Snijders, C., Buskens, V., & Cheshire, C. (eds), Trust and Reputation*. Russell Sage: New York. Forthcoming.
- Beckman, C. M., Haunschild, P. R., & Phillips, D. J. 2004. Friends or Strangers? Firm-Specific Uncertainty, Market Uncertainty, and Network Partner Selection. *Organization Science*, **15**(3), 259–275.
- Berg, Joyce, Dickhaut, John, & McCabe, Kevin A. 1995. Trust, Reciprocity and Social History. *Games and Economic Behavior*, **10**, 122–142.
- Boero, Riccardo, & Squazzoni, Flaminio. 2005. Does Empirical Embeddedness Matter? Methodological Issues on Agent-Based Models for Analytical Social Science. *Journal of Artificial Societies and Social Simulation*, **8**(4), 6.
- Boero, Riccardo, Bravo, Giangiacomo, Castellani, Marco, Laganà, Francesco, & Squazzoni, Flaminio. 2009a. Pillars of Trust: An Experimental Study on Reputation and Its Effects. *Sociological Research Online*, **14**(5), 5.
- Boero, Riccardo, Bravo, Giangiacomo, Castellani, Marco, & Squazzoni, Flaminio. 2009b. Reputational cues in repeated trust games. *Journal of Socio-Economics*, **38**(6), 871–877.
- Boero, Riccardo, Bravo, Giangiacomo, Castellani, Marco, & Squazzoni, Flaminio. 2010. Why Bother with What Others Tell You? An Experimental Data-Driven Agent-Based Model. *Journal of Artificial Societies and Social Simulation*, **13**(2), 6.
- Bravo, Giangiacomo. 2010. Agents' beliefs and the evolution of institutions for common-pool resource management. *Rationality and Society*, **22**(3), In press.
- Buchanan, Mark. 2002. *Nexus: small worlds and the groundbreaking science of networks*. New York: W.W. Norton.
- Buskens, V., & Weesie, J. 2000. Cooperation via Social Networks. *Analyse and Kritik*, **22**, 44–74.
- Buskens, V., Corten, R., & Weesie, J. 2008. Consent or Conflict: Coevolution of Coordination and Networks. *Journal of Peace Research*, **45**, 205–222.
- Buskens, Vincent, & Raub, Werner. 2008. Rational choice research on social dilemmas: embeddedness effects on trust. *In: Wittek, R., Snijders, T.A.B., & Nee, V. (eds), Handbook of Rational Choice Social Research*. New York: Russell Sage.
- Calvó-Armengol, A. 2001. On Bargaining Partner Selection When Communication Is Restricted. *International Journal of Game Theory*, **30**, 503–515.
- Camerer, Colin F. 2003. *Behavioral Game Theory. Experiments in Strategic Interaction*. New York / Princeton: Russell Sage Foundation / Princeton University Press.
- Centola, Damon, Willer, Robb, & Macy, Michael. 2005. The Emperor's Dilemma: A Computational Model of Self-Enforcing Norms. *American Journal of Sociology*, **110**(4), 1009–40.

- Chmura, T., & Pitz, T. 2007. An Extended Reinforcement Algorithm for Estimation of Human Behaviour in Experimental Congestion Games. *Journal of Artificial Societies and Social Simulation*, **10**(2), 1.
- Cohen, M., Riolo, R., & Axelrod, R. 2001. The Role of Social Structure in the Maintenance of Cooperative Regimes. *Rationality and Society*, **13**, 5–32.
- Coleman, James S. 1990. *Foundations of Social Theory*. Harvard: Harvard University Press.
- Corten, Rense, & Buskens, Vincent. 2009. *Co-evolution of Conventions and Networks: An Experimental Study*. ISCORE Paper 44: <http://www.uu.nl/uupublish/content/CortenBuskens_ISCORE.pdf>.
- Cronk, Lee. 2007. The influence of cultural framing on play in the trust game: a Maasai example. *Evolution and Human Behavior*, **28**, 352–358.
- Dal Forno, A., & Merlone, U. 2004. From Classroom Experiments to Computer Code. *Journal of Artificial Societies and Social Simulation*, **7**(2), 2.
- De Vos, H., Smaniotto, R., & Elsas, D. A. 2001. Reciprocal Altruism under Conditions of Partner Selection. *Rationality and Society*, **13**, 5–32.
- Deadman, Peter, Schlager, Edella, & Gimblett, Randy. 2000. Simulating Common Pool Resource Management Experiments with Adaptive Agents Employing Alternate Communication Routines. *Journal of Artificial Societies and Social Simulation*, **3**(2), 2.
- Dutta, B., Ghosal, S., & D., Ray. 2005. Farsighted Network Formation. *Journal of Economic Theory*, **122**, 143–164.
- Ebenhöh, Eva, & Pahl-Wostl, Claudia. 2008. Agent Behavior Between Maximization and Cooperation. *Rationality and Society*, **20**(3), 227–252.
- Eguíluz, Víctor M., Zimmermann, Martín G., Cela-Conde, Camilo J. Cela-Conde, & San Miguel, Maxi San Miguel. 2005. Cooperation and the Emergence of Role Differentiation in the Dynamics of Social Networks. *American Journal of Sociology*, **110**(4), 977–1008.
- Fischbacher, Urs. 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, **10**, 171–178.
- Flache, Andreas. 2001. Individual Risk Preferences and Collective Outcomes in the Evolution of Exchange Networks. *Rationality and Society*, **13**, 304–348.
- Gulati, R. 1995. Performance, Aspirations and Risky Organizational Change. *Administrative Science Quarterly*, **43**, 58–86.
- Hales, D. 2000. Cooperation without Space or Memory: Tags, Groups and the Prisoner's Dilemma. *Pages 157–166 of: Moss, S., & Davidsson, P. (eds), Multi-Agent Based Simulation*. Berlin-Heidelberg: Springer Verlag.
- Haruvy, E., Roth, A., & Ünver, M. 2006. The Dynamics of Law-Clerk Matching: An Experimental and Computational Investigation of Proposals for Reform of the Market. *Journal of Economic Dynamics and Control*, **30**(3), 457–486.

- Hauk, Esther. 2001. Leaving the Prison: Permitting Partner Choice and Refusal in Prisoner's Dilemma Games. *Computational Economics*, **18**, 65–87.
- Jackson, M. O., & Watts, A. 2002. The Evolution of Social and Economic Networks. *Journal of Economic Theory*, **106**, 265–295.
- Jackson, M. O., & Wolinsky, A. 1996. A Strategic Model of Social and Economic Networks. *Journal of Economic Theory*, **71**, 44–74.
- Jager, Wander, & Janssen, Marco A. 2003. Using Artificial Agents to Understand Laboratory Experiments of Common-Pool Resources with Real Agents. *Pages 75–102 of: Janssen, Marco A. (ed), Complexity and Ecosystem Management: The Theory and Practice of Multi-Agent Systems*. Cheltenham: Edward Elgar.
- Janssen, Marco A., & Ostrom, Elinor. 2006. Empirically Based, Agent-based models. *Ecology and Society*, **11**(2), 37.
- Joyce, David, Kennison, John, Densmore, Owen, Guerin, Steven, Barr, Shawn, Charles, Eric, & Thompson, Nicholas S. 2006. My Way or the Highway: a More Naturalistic Model of Altruism Tested in an Iterative Prisoners' Dilemma. *Journal of Artificial Societies and Social Simulation*, **9**, 4.
- Kagel, J., & Roth, A. 2000. The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment. *Quarterly Journal of Economics*, **115**, 201–235.
- Kahneman, Daniel, & Tversky, Amos (eds). 2000. *Choices, Values and Frames*. Cambridge: Cambridge University Press / Russell Sage Foundation.
- Keser, Claudia. 2003. Experimental Games for the Design of Reputation Management Systems. *IBM Systems Journal*, **42**, 498–506.
- King-Casas, Brooks, Tomlin, Damon, Anen, Cedric, Camerer, Colin F., & Montague, Steven R. Quartzand P. Read. 2005. Getting to Know You: Reputation and Trust in a Two-Person Economic Exchange. *Science*, **308**, 79–83.
- Knoch, Daria, Schneider, Frederic, Schunk, Daniel, Hohmann, Martin, & Fehr, Ernst. 2009. Disrupting the prefrontal cortex diminishes the human ability to build a good reputation. *Proceedings of the National Academy of Sciences USA*, **106**(49), 20895–20899.
- Kollock, Peter. 1994. The Emergence of Exchange Structures: An Experimental Study of Uncertainty, Commitment, and Trust. *American Journal of Sociology*, **100**(2), 313–345.
- Mark, N. 1998. Beyond Individual Differences: Social Differentiation from First Principles. *American Sociological Review*, **63**, 309–330.
- Molm, L. D., Takahashi, N., & Peterson, G. 2000. Risk and Trust in Social Exchange: An Experimental Test of a Classical Proposition. *American Journal of Sociology*, **105**(5), 1396–1427.
- North, Douglass C. 2005. *Understanding the Process of Economic Change*. Princeton, NJ: Princeton University Press.

- Ortmann, A., Fitzgerald, J., & Boeing, C. 2000. Trust, Reciprocity, and Social History: A Re-examination. *Experimental Economics*, **3**, 81–100.
- Pinyol, I., Paolucci, F., Sabater-Mir, J., & R., Conte. 2008. Beyond Accuracy. Reputation for Partner Selection with Lies and Retaliation. *Pages 128–140 of: Antunes, L., Paolucci, M., & Norling, E. (eds), Multi-Agent-Based Simulation VIII*. Berlin-Heidelberg: Springer Verlag.
- Podolny, J. 2001. Networks as the Pipes and Prisms of the Market. *American Journal of Sociology*, **107**, 33–60.
- Pujol, J. M., Flache, A., Delgado, J., & Sanguesa, R. 2005. How Can Social Networks Ever Become Complex? Modelling the Emergence of Complex Networks from Local Social Exchanges. *Journal of Artificial Societies and Social Simulation*, **8**(4), 12.
- R Development Core Team. 2009. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
- Rauhut, H., & Junker, M. 2009. Punishment Deters Crime Because Humans are Bounded in Their Strategic Decision-Making. *Journal of Artificial Societies and Social Simulation*, **12**(3), 1.
- Riolo, Rick L., Cohen, Michael D., & Axelrod, Robert. 2001. Evolution of Cooperation Without Reciprocity. *Nature*, **414**, 441–443.
- Rouchier, J. 2003. Re-Implementation of a Multi-Agent Model Aimed at Sustaining Experimental Economic Research: The Case of Simulations with Emerging Speculation. *Journal of Artificial Societies and Social Simulation*, **6**(4), 7.
- Slonim, R., & Garbarino, E. 2008. Increases in Trust and Altruism from Partner Selection: Experimental Evidence. *Experimental Economics*, **11**, 134–153.
- Watts, Duncan J. 1999. Networks, Dynamics and the Small-World Phenomenon. *American Journal of Sociology*, **105**(2), 493–526.
- Watts, Duncan J., & Strogatz, Steven H. 1998. Collective dynamics of ‘small-world’ networks. *Nature*, **393**, 440–442.
- White, Harrison C. 2004. *Markets from Networks: Socioeconomic Models of Production*. Princeton: Princeton University Press.
- Willer, Robb, Kuwabara, Ko, & Macy, Michael W. 2009. The False Enforcement of Unpopular Norms. *American Journal of Sociology*, **115**(2), 451–490.
- Yagamashi, T., Cook, K. S., & Watabe, M. 1998. Uncertainty, Trust, and Commitment Formation in the United States and Japan. *American Journal of Sociology*, **104**(1), 165–194.